

Dictionary-Based Instruction Compression and the TMS320C6x Architecture

Andrew Aarestad

Joseph Thomas

Kyaw Min

1. Introduction

As solid-state memory technology develops and becomes cheaper, design parameters for devices that use it can change as well. For years, one of the primary goals of computer engineering, especially in the embedded systems field, was the reduction of the memory footprint requirements of software. Now that memory is becoming increasingly larger and more affordable, the design focus is shifting away from footprint reduction to reduction of power consumption. Power consumption has always been a concern for computer engineers, but only recently has it moved to the forefront of engineering design metrics. New methods of data storage are being developed that not only reduce memory footprint requirements but also reduce memory bus traffic, a major contributor to the overall system power consumption[REF]. Instruction compression is one of the techniques being applied to reduce power consumption.

In this paper, we present several schemes that we have developed for use in the Texas Instruments TMS320C6x VLIW architecture. We explore the various performance tradeoffs of modifications to the original architecture, adding compression to the scheme in several places [1].

2. Design Parameters

Energy Savings

The primary design goal of the compression system is to reduce the energy requirements of the system. This is achieved by increasing the cache hit rate. Fewer cache misses means the CPU has to go off-chip to satisfy its instruction requests less. However, we do not provide an analysis of the gains in terms of power savings since the power requirements of the memory system varies by implementation. Instead, we provide an analysis of the hit rate gains which, given information about the power consumption of the memory bus, could be converted to data about the power consumption.

Transparency

One of the design goals for this project is transparency to the CPU core. This means that one of our requirements is that modifications to the core are unnecessary and unwanted. Our modifications include positioning a decompression unit between different levels of instruction cache. This does mean modification of the chip since the cache is on-chip, but does not require modification of the core. The idea is to allow the core to behave exactly as if there were no modifications. This is an important feature when designing for processors with proprietary cores such as the one we are using [REF].

Cache Bus Widths

The original system as defined in TI's documentation calls for L1 and L2 caches line sizes of 512 and 1024 bits, respectively [2]. The TMS is a VLIW processor, which means that it processes a fetch packet of up to eight instructions per clock cycle. Each instruction is 32 bits, so the cache lines for the L1 and L2 caches represent 2 and 4 fetch packets, respectively. To implement the compression schemes we have proposed, redesign of the cache busses would be required since fewer bits would need to be transferred to satisfy the same requests.

Chip Cost

We have attempted to design our compression schemes in such a way that the cost of the chip does not increase. Since we have allocated cache-speed memory for use in decompression, we have 'borrowed' the memory from the L2 cache. Thus, in our analysis we have reduced the size of the L2 cache to keep the total memory requirements of the chip constant.

3. Compression Schemes

Both of these modifications of the TI TMS320C6x architecture compress machine instructions by placing an instruction decompression unit on the chip (Fig 1). Instructions are compressed before they are placed in main memory from 32 to 16 bits. The compressed version of the program and a decompression dictionary are loaded into the DSP's memory. The decompression unit has a small amount of local memory that is used to hold the dictionary. When the CPU requests a 256-bit fetch packet located at a certain address, the decompression unit intercepts the request, translates the address into the compressed address space, fetches the compressed instructions, and decompresses them before returning them to the CPU. The difference between the two schemes we propose here is the location of the decompression unit in the cache system.

The following descriptions of what happens during an instruction fetch illustrate our changes to the original system and give a side-by-side comparison of the actions of both our schemes.

Scheme 1

1. Address is generated in CPU core
2. Address is sent to L1 cache. As in the original configuration, the L1 cache contains uncompressed fetch packets. If the requested fetch packet is not located in the L1 cache, the cache updates.




Figure 1 - Shows the placement of the decompression unit for the schemes.

3. Fetch packet is returned if there is a L1 cache hit, otherwise the address is sent to the decompression unit.
4. Address is translated into compressed address space. Since the compressed instructions are half the size of the original, the address of a compressed instruction will be half of the original request.
5. Address is sent to the L2 cache.
6. The L2 cache is searched for the compressed address.
7. Compressed cache line is returned to decompression unit if there is a L2 hit, otherwise the requested address is sent to the DMA unit.
8. Requested instructions are fetched from main memory.
9. L2 cache line is received from main memory, L2 cache updated.
10. Compressed cache line is sent to decompression unit.
11. Cache line is decompressed.
12. Decompressed cache line is sent to the L1 cache.
13. L1 cache is updated.
14. Decompressed fetch packet is sent to CPU.

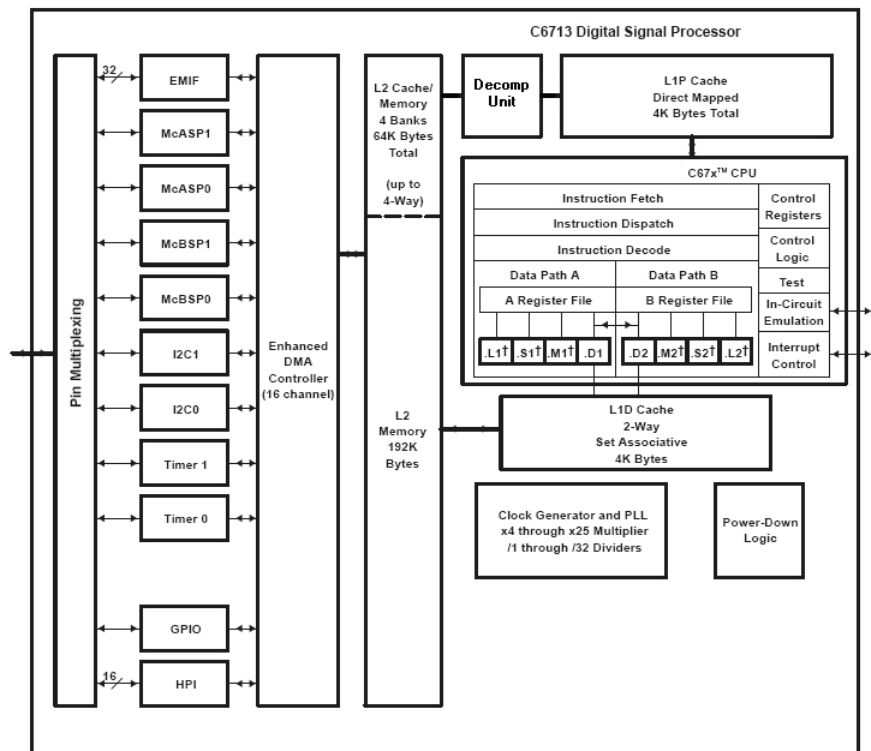


Figure 2 - The TMSC6000 architecture with decompression unit inserted for scheme 1.

Scheme 2

1. Address is generated in CPU core
2. Address is sent to decompression unit
3. Address is translated into compressed address space
4. Compressed address is sent to L1 cache
5. L1 is searched, if it misses, the address is sent to the L2 cache
6. L2 is searched, if it misses, it is updated from main memory
7. Instructions are fetched from main memory
8. Instructions are received, L2 frame is updated
9. Compressed cache line is sent to L1
10. L1 cache line is updated
11. Compressed fetch packet is sent to decompression unit
12. Fetch packet is decompressed
13. Decompressed fetch packet is sent to the CPU

4. Performance Analysis

We have evaluated the performance of our compression schemes using the following metrics:

- Memory Footprint Reduction
- Change in Cache Hit Rate
- Instruction Access Time Penalty

Memory Footprint

Both schemes use static-length 16-bit codewords. This gives us an even 50% reduction of the program size, but the file must also contain the decompression dictionary. In smaller programs, this may actually result in code expansion. See Fig. 2 for data relating to memory footprint performance of our schemes.

Cache Hit Rate

The primary goal of our compression was to increase the effective cache hit rate. By compressing the caches, we are able to store more fetch packets in cache for the same cache size. With more chances for a cache hit, there will be fewer cache misses.

Access Time Penalty

The tradeoff that comes with the cache hit rate increase is a penalty in instruction access time. A compressed fetch packet is half as long as an uncompressed one. Because the memory controller must bring the instructions of the fetch packet into the chip in serial, only half as many of these fetches would occur per fetch packet, reducing the L2 miss penalty significantly. However, the penalty is increased by the delay caused by decompression.

5. Conclusions

To analyze the performance of our schemes, we simulated several DSP-specific applications on Vinodh Cuppu's c6xsim cycle-accurate simulator [4]. We then ran the instruction trace generated through a cache performance simulator that we developed. This simulator is more thoroughly documented in our document "TMSCache Cache Performance Simulator User's Guide." [REF]

Scheme 1 has the same L1 cache performance as the original system since it leaves the L1 cache uncompressed. Gains are noticed in the L2 cache hit rate, see Fig. 3 for data. Since the L2 cache is updated serially, the L2 penalty is reduced by compression. This penalty reduction compensates for the penalty increase caused by decompression and produces an overall decrease in average instruction access time. See Fig. 4 for data about the average access times.




Figure 3

Scheme 2 had better cache miss rates than Scheme 1 or the original, but does so at the cost of access time. Since it must decompress every instruction issued, the decompression penalty adds up fast. As is illustrated in Fig. 3, the average access time is increased by about a factor of three. For a non-time-critical system in which energy savings is paramount, a system such as this could be desirable, but for most systems this is too large of a penalty.




Figure 4

In conclusion, our data shows that implementing a compression technique such as the ones we illustrated here can result in both energy savings and performance gains without increasing the cost of the chip.

References

- [1] Luca Benini, Alberto Macii, and Enrico Macii, “Minimizing Memory Access Energy in Embedded Systems by Selective Instruction Compression,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 10, No. 5, pp. 521-530, Oct 2002.
- [2] TMS320C6713 Floating-Point Digital Signal Processor: Texas-Instruments, 2005.
- [3] TMS320C6000 CPU and Instruction Set Reference Guide: Texas-Instruments.
- [4] Vinodh Cuppu, Cycle Accurate Simulator for TMS320C62x, 8 way VLIW DSP Processor: University of Maryland, College Park, 1999.
- [5] Luca Benini, Francesco Menichelli, and Mauro Olivieri, “A Class of Code Compression Schemes for Reducing Power Consumption in Embedded Microprocessor Systems,” *IEEE Transactions on Computers*, vol. 53, No. 4, pp. 467-482. Apr. 2004.
- [6] Montserrat Ros, and Peter Sutton. “Code Compression Based on Operand-Factorization for VLIW Processors,” *IEEE Proceedings of the Data Compression Conference*, pp. 1-5, Apr 14, 2004.
- [7] J. L. Hennessy and D.A. Patterson, *Computer Architecture – A Quantitative Approach*, 3rd ed. San Mateo, CA: Morgan Kaufmann, 1996.